User-driven Optimization of Shared Autonomy in Assistive Robotics

Deepak Gopinath[1,2], Siddarth Jain[1,2] and Brenna D. Argall[1,2]
[1]Northwestern University, Evanston, Illinois, 60208, USA
[2]Rehabilitation Institute of Chicago, Chicago IL, 60211 USA

I. Introduction and Motivation

For people with severe motor impairments assistive and rehabilitation machines such as assistive robotic arms and upper or lower limb prostheses are crucial in reducing their dependence on caretakers and increasing the ability to perform activities of daily life. However, for many, the control of such devices remains a challenge—for example, due to their physical impairments or limitations of the control interfaces. The introduction of partial autonomy to these devices—in which the control is shared between the human and robotics autonomy—aims to help reduce the cognitive and physical burden on the user. The reduced bandwidth of the control signals generated by motor-impaired users makes them more reliant on the interaction with the autonomy, and also less adaptable and more vulnerable to any arbitrariness present in the system—for example, the choice of control interfaces and mappings, or the exact specification of how control is shared between the user and the autonomy.

Since users differ in their physical abilities and desired amount of assistance, *customization* of the amount of assistance is critical for the adoption of assistive shared-control systems. Predefined assistance levels may not remain optimal for the user in the long term, as the need for assistance may increase or decrease. One way to accomplish customization is to tune the system parameters which will bring about a change in the human-robot interaction and the final behavior. The aim is to optimize the human-robot interaction during task performance. A simple choice of optimality criterion is to consider task-related performance metrics such as minimizing the time taken and energy expended.

Our insight is that if we entrust the task of customization to the users, they will tune the system in such a way that likely achieves their preference, and perhaps also improves satisfaction and performance. This personal optimality criterion may not be a simple time–optimal or energy–optimal objective function. Moreover, the needs of each user can undergo variation due to factors such as fatigue, loss of ability due to degenerative disease or improvement due to rehabilitation.

Many optimization techniques have been adopted to generate different strategies for control sharing; for example, formulating the problem as a POMDP and inferring a distribution over goals from user commands via Maximum Entropy Inverse Optimal Control [2], or concatenating energy-optimal motion primitives to create optimal trajectories [3]. Although these approaches result in improved task performance (task completion time, control effort), the assistance schemes are mixed in terms of user acceptance. In particular, there are instances of assistance resulting in higher user dissatisfaction [2], and users preferring to be in control and more cautious [3]. In other studies users find the assistance at times to be uncooperative and are willing to tolerate a loss of control only when there is a significant improvement in performance [4].

The need for high user satisfaction is crucial for the acceptance of robot autonomy by the end-users in the assistive domain. This notion motivates our approach to engage the end-user in the customization procedure under the assumption that the user knows what is preferred by him/herself.

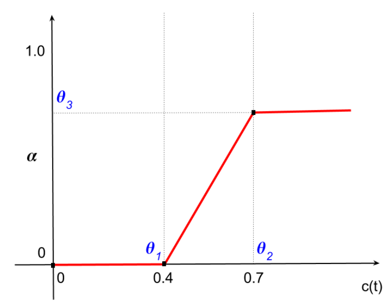

Fig. 1. A prototypical arbitration function, parameterized by $\boldsymbol{\theta} = \{\theta_1, \theta_2, \theta_3\}$.

To ground our formalism, we present a first implementation in which the reasoning between the user control and the robot policy is a function of the system's confidence in its inference of human intent, with tunable parameters (Figure 1) [1]. The parameters ($\boldsymbol{\theta} = \{\theta_1, \theta_2, \theta_3\}$) affect the onset of assistance, aggressiveness and the maximum amount of assistance offered by the robot. In the prototype user-driven optimization system developed for this work verbal commands from the user are mapped to changes in the parameters by the system operator. We also present results from a pilot study which shows that spinal cord injury (SCI) subjects were able to achieve task performance comparable to that of uninjured subjects with customization. Notably, the amount of assistance was not always optimized for task performance. Some subjects favored retaining more control during the execution over better task performance.

II. Formalism and Optimization Procedure

Let $\boldsymbol{x}(t)$ denote the state of the system at time t. Let $\boldsymbol{\theta}(t)$ be the set of tunable parameters that will affect the amount of control shared between the human and the robot. The other control inputs to the system are $\boldsymbol{u}_h(t)$ and $\boldsymbol{u}_r(t)$, the control

Verbal Cue	Parameters Changed	Amount of change
"More"	$\theta_3 \uparrow$, $\theta_2 \downarrow$, $\theta_1 \downarrow$	$\delta\theta \leftarrow \delta\theta$
"Less"	$\theta_3 \downarrow$, $\theta_2 \uparrow$, $\theta_1 \downarrow$	$\delta\theta \leftarrow \delta\theta$
"Little More"	$\theta_3 \uparrow$, $\theta_2 \downarrow$, $\theta_1 \uparrow$	$\delta\theta \leftarrow \frac{1}{2}\delta\theta$
"Little Less"	$\theta_3 \downarrow$, $\theta_2 \uparrow$	$\delta\theta \leftarrow \frac{1}{2}\delta\theta$

TABLE I
MAPPINGS FROM VERBAL CUES TO PARAMETERS CHANGED
($\uparrow$ indicates a positive $\delta\theta$ and $\downarrow$ denotes a negative $\delta\theta$)

Control Mappings		
Mode	**3D**	**2D**
1	v_x, v_y, v_z	v_x, v_y
2	$\omega_x, \omega_y, \omega_z$	v_x, v_z
3	—	ω_x, ω_y
4	—	ω_z

TABLE II
OPERATIONAL PARADIGMS FOR THE TELEOPERATION INTERFACE

commands generated respectively by the user and autonomous robot policy at time t.

The control signal from the robot autonomy is generated by a function $f(\cdot) \in \mathcal{F}$,

$$\boldsymbol{u}_r(t) \leftarrow f(\boldsymbol{x}(t)) \tag{1}$$

where $\mathcal{F}$ is the set of all control behaviors corresponding to different tasks.

We assume that the control command $\boldsymbol{u}_h(t)$ is generated by a function of $g(\cdot) \in \mathcal{G}$,

$$\boldsymbol{u}_h(t) \leftarrow g(\boldsymbol{x}(t)) \tag{2}$$

where $\mathcal{G}$ is the set of user behaviors corresponding to different tasks. $g(\cdot)$ is simply a symbolic representation of the mapping function that generates $\boldsymbol{u}_h(t)$ and is completely unknown to the autonomous system.

The shared control system makes use of function $\boldsymbol{\beta}(\cdot)$, parameterized by $\boldsymbol{\theta}$

$$\boldsymbol{u}(t) \leftarrow \boldsymbol{\beta}_\theta(\boldsymbol{u}_h(t), \boldsymbol{u}_r(t)) \tag{3}$$

which arbitrates between the control commands from the user and the robot policy to produce control command $\boldsymbol{u}(t)$ executed by the robot.

A key insight in our formulation is that, for a time–varying function $\boldsymbol{\beta}(\cdot)$, the parameters themselves can be functions of time and therefore may be interpreted as control signals. Then the dynamics of the system can be written as

$$\dot{\boldsymbol{x}}(t) \leftarrow \boldsymbol{a}(\boldsymbol{x}(t), \boldsymbol{\theta}(t), \boldsymbol{u}_h(t), \boldsymbol{u}_r(t), t) \tag{4}$$

where $\boldsymbol{a}(\cdot)$ is in general a non-linear, time–varying function. The problem of finding the set of parameters $\boldsymbol{\theta}(t)$ that will generate the optimal human-robot interaction and task performance thus may be formulated as an optimization problem.

The elements of the framework $f(\cdot)$, $g(\cdot)$, $\beta(\cdot)$ and $a(\cdot)$ are system-specific, and different choices of these functions will have drastically different impact on task performance and user satisfaction. As discussed in the introduction, the impact is anticipated to be all the greater on motor-impaired subjects.

A. Optimization

Typically optimization is performed over all control signals that are inputs to the system. In our system, however, the control commands from the human and the robot are treated as given quantities, and the goal rather is to optimize the interaction parameters $\boldsymbol{\theta}(t)$. Therefore, optimization is performed only with respect to a subset ($\boldsymbol{\theta}(t)$) of the entire control space.

Since we do not want to reduce the assistive capabilities of our system, and we have a human in the loop, our insight is that the optimization task can be performed by the user him/herself. However, there may be a variety of unmeasurable factors influencing the cost function, and determining the exact mathematical form for the cost function may be an intractable problem. Making any kind of approximation to simplify the cost function in turn will affect the robustness and efficacy of the assistive system.

B. User-Driven Optimization of the Arbitration Parameters

In this first exploration of our interactive optimization procedure, verbal commands from the human subject are translated to changes in $\boldsymbol{\theta}$ by the system operator.

A change in assistance level can be achieved by modulating one or more of the $\theta_i \in \boldsymbol{\theta}$, according to $\theta_i = \theta_i \pm \delta\theta_i$. In our implementation, at initialization $\delta\theta_i = 0.1$. The value of $\delta\theta_i$ is adaptive, and is halved if a request to increase assistance is immediately followed by a request to decrease and vice versa (in order to avoid oscillatory behavior). This procedure is analogous to a three dimensional gradient descent algorithm in the sense that both the direction and the magnitude of change are being updated at every optimization step. Table I provides the mappings between common verbal cues, the parameters changed and the values of $\delta\theta$. We chose to modulate more than one parameter at a time as it helped to make the change in assistance level more perceivable to the user.

III. EXPERIMENTAL SETUP

The experiments were performed using the MICO robotic arm (Kinova Robotics, Canada) which is specifically designed for assistive purposes.

A. Control Interface

The human control command $\boldsymbol{u}_h(t)$ is captured via a teleoperation interface, that consists of a 3-axis joystick operated under two different mapping paradigms (namely, *3D* and *2D*) (Tbl II). The joystick signals are mapped to the translational (v_x, v_y, v_z) and rotational ($\omega_x, \omega_y, \omega_z$) velocities of the end-effector in Cartesian space.

B. Task Descriptions

Three tasks were developed for our pilot study.

Simple Reaching (R): The user teleoperates the robotic arm to reach a coffee carafe placed in front of the robotic arm.

Reaching for Grasping (RfG): The user teleoperates the robotic arm to reach one of two objects on the table with a pose suitable for grasping, as the robot arm provides assistance. There is a near object (mug) and a far object (box), each of which requires a different orientation of the gripper for grasping (side and top, respectively) and accordingly also different approach trajectories during reaching.

Reaching for Scooping (RfS): The user teleoperates the robotic arm to reach for one of two objects on the table with a pose suitable for a scooping motion, as the robot arm provides assistance. There is a near object and a far object (both bowls), each of which requires a different approach trajectory. For this task, the end effector of the robotic arm is fitted with a spoon which must be inserted into the bowl.

C. Methods

Subjects: For this exploratory study 17 subjects were recruited—13 uninjured control subjects (mean age: 26 ± 4, 8 males and 5 females) and 4 SCI subjects (mean age: 35 ± 14, all males, C3-C5 injury levels). Ten subjects (7 uninjured and 3 SCI) used the *3D* interface paradigm, and the remaining subjects used the *2D* paradigm.

Protocol: Each user performed all three tasks. The purpose of task *R* was to get the user accustomed to the control interface and to the different assistance levels. Data was then collected on the remaining two tasks (*RfG*, *RfS*). For the *RfG* and *RfS* trials, the user first operated the system in full teleoperation mode (*tel*) and then under three predefined assistance levels (*min, mid* and *max*). After this phase, the subject was given the option to further customize the assistance level, resulting in assistance level *custom*. Three trials were collected for *min, max* and *custom* assistance levels. For the first (non-practice) task, the baseline from which customization began was the *mid* level assistance, with level *custom* being the result after customization. For the second task, customization began at this level *custom*, with the option to further customize resulting in level *custom* for the second task.

Metrics: *Task Completion time* is the amount of time spent accomplishing a task. *Mode Switches* (an indirect measure of effort) refers to the number of times the subject switched between the various modes of the control interface.

IV. Results

From our pilot study we saw that user-driven customization in general improves task performance and helps to reduce performance differences between uninjured and SCI subjects. In Figure 2, the difference between task completion times for uninjured and SCI subjects drops steadily from *tel* to *custom* assistance levels. That is, with increased assistance, the performance of SCI subjects was comparable to that of uninjured subjects and this was *maximized* under customization. The variance in the data also *decreases* with customized assistance, showing the performance to become more consistent.

For the *custom* assistance type, even though the task completion times are comparable to those of the *max* assistance type, the number of mode switches is greater than those of *max*. Thus, it is not uniformly the case that mode switches are minimized from low to high assistance. This observation provides insight that the true cost function that the user is optimizing is more complex that a simple time-optimal or energy-optimal cost function.

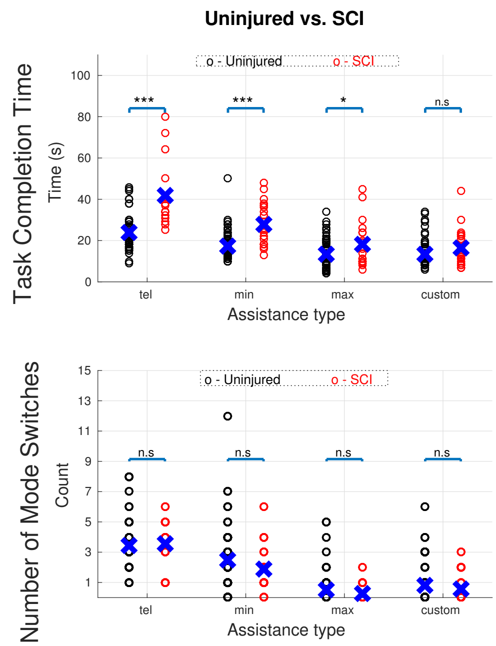

Fig. 2. Task completion time (top row) and number of mode switches (bottom row) for uninjured vs. SCI subjects. (***) $p < 0.001$, (**) $p < 0.01$ and (*) $p < 0.05$, (n.s) not significant.

V. Conclusion

In this work, we have introduced a system for user-driven customization which is presented as a constrained non-linear optimization problem. Unlike standard optimization problems in which the form of the cost function is predetermined in this work no such assumptions were made. Instead, the end user was allowed to directly perform the optimization procedure. The aim is that this will lead to higher user satisfaction. An interactive user-driven customization system was developed to ground the formalism and selected results from the pilot study were presented.

Acknowledgments

Many thanks to Jessica Presperin Pedersen, OTR/L, ATP/SMS, for recruiting the SCI subjects, and to Samuel Schlesinger for assistance during the subject studies. Research reported in this publication was supported by the NIBIB & NICHD under award number R01EB019335.

References

[1] Anca Dragan and Siddhartha Srinivasa. Formalizing assistive teleoperation. In *Proceedings of RSS*, 2012.

[2] S. Javdani, S. Srinivasa, and A. Bagnell. Shared autonomy via hindsight optimization. In *Proceedings of RSS*, 2015.

[3] Martin Lawitzky, Melanie Kimmel, Peter Ritzer, and Sandra Hirche. Trajectory generation under the least action principle for physical human-robot cooperation. In *Proceedings of the ICRA*, 2013.

[4] Erkang You and Kris Hauser. Assisted teleoperation strategies for aggressively controlling a robot arm with 2D input. In *Proceedings of RSS*, 2012.